


Alternating-Z Tripler with High Dynamic Range

Dave Eimerl
David Milam

October 10, 1996



DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This report has been reproduced
directly from the best available copy.

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information
P.O. Box 62, Oak Ridge, TN 37831
Prices available from (615) 576-8401, FTS 626-8401

Available to the public from the
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Rd.,
Springfield, VA 22161

to: distribution

From: David Eimerl and David Milam

Subject: Alternating-Z tripler with high dynamic range.

I. Summary.

1. A KD*P alternating-Z tripler suitable for NIF consists of two detuned Type I doublers and one aligned Type II mixer. The dynamic range for efficient tripling is predicted to be larger for this arrangement than for a standard 2-crystal tripler.

2. We tested an alternating-Z tripler that had doublers with thicknesses of 13 and 10 mm, and a 10-mm-thick mixer. All of the crystals were 80% deuterated.

3. The model for this tripler provided an excellent description of the efficiencies that were measured using pulses with duration of 1 or 6 ns.

4. We confirmed the importance for alternating-Z triplers of controlling the phase differences arising from wavelength dispersion in the air gap between the doublers.

5. The confirmed wide dynamic range of this arrangement is of interest for conversion of the 21-ns Haan pulse, or for conversion of beams with significant intensity fluctuations. Compensation for wavelength dispersion in the gap between the crystals is not an issue if the crystals are in vacuum.

II. Introduction.

Frequency conversion of intense laser light by harmonic generation in a nonlinear optical material was first reported many years ago and is now a well-known element of laser-electrooptic technology [1,2]. In applications, a laser beam at one frequency (e.g. 1053 nm) is directed through carefully oriented crystals of a nonlinear material, that respond nonlinearly and generate a laser beam at the output wavelength

with very similar propagation parameters, with less than 100% efficiency. One approach to improving the conversion efficiency in practice is to use more than the minimum number of crystals. The quadrature multiple crystal scheme for frequency doubling has already been discussed[3]. In this paper we introduce new multiple crystal approach for frequency tripling, the modified alternating-Z design, and present data documenting its performance for tripling 1053 nm to 352 nm. Note that in this paper we are interested in large aperture beams; the discussion may not apply to focused beams or other optical arrangements.

The value of multiple crystal conversion schemes is that they permit significant increases in the range of intensity over which the conversion efficiency is high, and are therefore more robust designs. The key parameter is the dynamic range. The theoretical maximum efficiency for second harmonic and third harmonic generation is controlled by the fundamental physics of the conversion process, as well as engineering factors such as optical losses at the crystal surfaces. Because the process is nonlinear and usually local and effectively instantaneous, the efficiency depends on the spatial and temporal intensity profile of the incident laser beam. The dynamic range is a measure of the intensity variation that can be tolerated (for a given wave-front shape) without significant loss in conversion efficiency. It is the ratio of the highest intensity for which the efficiency exceeds some desired value to the lowest intensity for which it exceeds the same value. Remarkably if the desired efficiency is not too close to the maximum efficiency the dynamic range is relatively insensitive to the desired efficiency but depends mostly on the wavefront flatness, laser power and the angular sensitivity of the crystal. For standard single-crystal schemes, the dynamic range Ω for doubling is at most 10-15, and for tripling it is about 3.

The dynamic range is a pure number, so defined to respect a scaling law for the conversion efficiency as the beam area is changed. The area scaling law for the maximum efficiency applies to the use of paraxial optics to change the beam area. If the incident beam is changed by a factor λ in linear dimension, and the crystal thickness also scaled by the same factor, then the conversion efficiency is unchanged. This was demonstrated in [2,3]. This scaling law is exact if linear optical absorption and other nonlinear effects such as two-photon absorption are neglected. In the presence of linear absorption the efficiency decreases as the beam area is increased, whereas for nonlinear effects the efficiency increases with beam area. Thus the conversion efficiency depends only on the laser power, rather than the power density. Because of the area scaling law, the design process fixes length ratios such as the ratio of crystal thickness to beam aperture, rather than absolute magnitudes. In particular while either the highest or the

lowest of the intensities that define the useful operating range in intensity can be varied at will by paraxial scaling, they both scale in the same way, so their ratio is constant. Therefore the dynamic range characterizing the robustness of the convertor is defined as a ratio. In comparing multiple crystal arrangements with standard large aperture designs the dynamic range so defined is the key figure of merit.

In general for any frequency conversion scheme there are only two parameters for each crystal, orientation, and thickness. Once these have been optimized to maximize the efficiency, the actual maximum reached depends only on the laser beam parameters and the crystal material constants. The area scaling law shows that the optimum can depend only on the laser power, wave-front quality and beam profiles, and not on the intensity itself.

In practice the optimum may be found by a design process as follows. Given a beam with specified wavefront and intensity profile, the convertor is designed in conjunction with an optical system that collimates the beam and enlarges its aperture to avoid optical damage to the crystals. Then the crystal lengths and orientations are determined to optimize the conversion efficiency. First the crystals are oriented to optimize the conversion efficiency for a low intensity beam with the same wave-front and profile shape. For some schemes, this is followed by another angular adjustment to off-set the crystals from this optimal orientation. Then the crystal lengths are selected to optimize the conversion efficiency for the full intensity beam. If necessary, the orientation is then finely adjusted at high intensity. Alternatively, instead of adjusting the crystal length an equivalent procedure is to adjust the beam aperture, thereby changing both the intensity on the crystal and the beam divergence. The difference between these two design procedures is slight.

Using this procedure or an equivalent one the optimum efficiency can be reached, and the crystal lengths and the optimum efficiency will depend only on the laser power, wave-front quality and beam profiles. The most significant parameter controlling the conversion efficiency is the wave-front quality of the incident beam, followed by the intensity profile.

The maximum efficiency for a given beam can be predicted with numerical models to better than 1% accuracy. Using these models frequency convertors can be designed reliably if the incident beam is known. In general, as the beam quality is changed the highest conversion efficiency will also change - the optimum conversion efficiency and crystal parameters will be best for beams with the flattest wave-fronts and the flattest intensity profiles, and will be reached by a procedure equivalent to that just described.

In practice it is difficult to know the incident beam parameters well enough to predict performance, and selecting the crystal lengths must proceed with incomplete information. The most robust conversion schemes are those with the highest dynamic range and therefore the least sensitivity to the beam quality. By increasing the dynamic range the conversion efficiency where the local, instantaneous intensity is low, can be made significantly larger. Thus for beams with strong, or unpredictable, local intensity variations, high dynamic range is highly desirable. Multiple crystal schemes offer significantly higher dynamic range over the standard designs. For example quadrature schemes for frequency doubling can increase the dynamic range from about 15 to several hundred. The alternating-Z schemes decrease the angular sensitivity and can preserve the dynamic range under conditions where the beam quality is worse, or where there are temperature gradients in the crystals.

While quadrature schemes increase the robustness dramatically for frequency doubling, a corresponding scheme for tripling has been lacking. In this paper we describe a high dynamic range tripling scheme based on a modification of the alternating-Z doubling scheme. The alternating-Z design was first proposed in the context of high average power frequency conversion, and has greater tolerance for temperature gradients and stress-optic effects than a single crystal doubler. In the tripling application this design is modified with the addition of angular offsets. We also present detailed experimental confirmation of the robust performance of the alternating-Z tripler.

III. High Dynamic Range Tripler.

The standard configuration [4] for efficient frequency tripling consists of a doubler that converts part of the incident light (1053 nm) to the second harmonic (527 nm) followed by a crystal that mixes the so-generated second harmonic with the remaining fundamental to generate the third harmonic at 351 nm. Details of this design and others can be found in Craxton [4] and Eimerl [3]. For optimum operation of the mixer crystal the ratio of incident second harmonic to first harmonic should be exactly two. A mix ratio $M=2$ can be achieved by the doubler at a particular input intensity in various ways, such as selecting the doubler length, or giving it an angular offset, or by rotating the polarization of the incident 1053 nm light. However, its dynamic range is another matter - the dynamic range of a single crystal doubler designed to operate at 67% conversion efficiency is small, $\Omega < 3$. Because the dynamic range of the tripling combination is constrained by the dynamic ranges of both the doubling and the mixing processes, the doubler strongly constrains the dynamic range of the combination to

values in this range (< 3). This occurs despite the fact that the mixer dynamic range for a single crystal convertor/mixer operating close to saturation (100% efficiency) can be in the range $5 < \Omega < 15$.

The dynamic range of doubling at 67% efficiency can be significantly greater using a two-crystal alternating-Z design. The basic idea of the alternating-Z approach is that the phase mismatch developed in the first crystal is compensated in the second crystal, and the basic idea of multiple crystal schemes is that crystals with different lengths perform most efficiently in different intensity ranges. We would therefore expect to have two crystals of different lengths, each with an angular offset to provide the correct mix ratio, and that the angular offsets would be such that the dephasing accumulated in the first would be compensated in the second. Therefore we would expect that the ratio of the detuning angles should be the inverse of the ratio of the crystal lengths, and should be of opposite sign.

This elementary picture is confirmed by detailed modeling of a pertinent problem, frequency tripling at the National Ignition Facility, where KDP (potassium dihydrogen phosphate) crystals will be used to convert 1053 nm beams 20-40 cm in aperture, having excellent wavefront quality. High dynamic range is highly desirable because the output pulse shape requirements at 351 nm involve a large intensity range 0.05 - 5 GW/cm². One problem with optimizing multiple crystal schemes is the potentially large number of independent variables. Each crystal has two parameters, a length and an angular offset, and can perhaps be of more than one type. The incident intensity must be varied to assess the dynamic range. The number of variables is then at least one more than twice the number of crystals, (7 for three-crystal devices) and is even greater if other incident beam parameters are included such as depolarization. Using Runge-Kutta codes we explored the parameter space, and confirmed that the dynamic range can be increased significantly using the above physical picture. While it is highly unlikely that other high-dynamic-range configurations exist involving two doublers and a single mixer, the parameter space is too large to rule this out absolutely. Using the physical model, approximate values for the variables were determined, and then the design was optimized for beams with an intensity variation and wave-front quality similar to that expected for the National Ignition Facility.

Several almost equivalent designs were found. The doublers can be either type I or type II, and the orientation of the doublers can be either tandem or alternating-Z. In tandem schemes the z-axes are parallel, whereas for alternating-Z schemes some crystals are rotated from the parallel case by 180° about the beam direction. We selected type I doublers, because these have smaller crystal boule size requirements than type II

crystals. However, there are several (almost) equally good configurations involving other crystal cuts. The incident beam was assumed to be 50 μrad in divergence. The beam profile was assumed to deliver most of the beam energy between 0.1 and 5 GW/cm^2 . With these guidelines the highest dynamic range was found for the configuration described in Table 1. A configuration very similar to this was set up and tested experimentally. The improved robustness of its performance was confirmed.

Note that it is also possible to add a second mixing crystal to increase the mixing dynamic range, this has only a slight effect on the performance because the dynamic range is still controlled by the doubler pair.

Type I/Type I/Type II, Alternating-Z	Optimum
Beam polarization	90°
Beam divergence	30-50
Intensity minimum	0.01
Intensity maximum	6
First doubler length	13
First doubler offset	400
Second doubler length	10
Second doubler offset	-520
Air path phase error	0
Mixer length	13
Mixer offset	0
Crystal deuteration level	80%

Table 1.
An optimized design for
an alternating-Z tripler.

IV. The Sign of the Nonlinear Coefficients.

Besides the standard tandem and alternating-Z configurations described above, cases where some z-axes are exactly reversed from these standard cases can also exist. Z-reversal can be a problem because the effective d-coefficient can change sign under Z-reversal, and care must be taken to ensure the correct sign in practical situations. Another factor affecting the sign of d_{eff} is the XY-cut of the crystal. The crystal has tetragonal symmetry, so that equivalent positions are found by rotating the crystal axes by $\pi/2$ about the z-axis. The d_{eff} changes sign under this rotation, so that there are two types of crystals. For example, for type I crystals, the input face may contain either a

crystalline X-axis, or a Y-axis. (For type II the input face may contain either an (X+Y) or an (X-Y) direction.) These two cuts will perform equally well as single crystal doublers but their d-coefficients have opposite signs.

The relative sign is critical for tandem and alternating-Z configurations because the second harmonic waves generated in the two crystals are coherent, and their relative phase must be controlled. The simplest way to see this is to consider a single crystal convertor of thickness L that has been cut into two crystals of equal length $L/2$. Clearly the pair together will perform as the original crystal if placed next to each other. However, if one of them were rotated so that its d-coefficient was of the opposite sign, this would cause the conversion to drop dramatically, because the second crystal would back-convert the light generated in the first one. Because the relative sign is so critical, it is essential to know the XY-cut as well as the positive z-axis direction for the crystals in order to obtain the desired performance.

In this connection we mention that both Summers et al [5] and Velsko et al [6] have exploited this coherence effect to measure the wavelength dispersion of the refractive index of gases. Velsko set up a pair of crystals in tandem with a tube of gas under pressure placed between them. The doubling efficiency (at low intensity) was measured as the pressure of the gas was varied. Summers measured the conversion efficiency at low intensity as a function of the spacing of two doublers that were in air at atmospheric pressure. In both experiments, a sinusoidal variation was observed, which arises from the wavelength dispersion of the gas.

The interpretation of this experiment is as follows. The second harmonic generated in any crystal is phase-locked to the incident fundamental. In this case phase-locking should be interpreted in terms of an instantaneous picture of the spatial variation of the electric fields of the fundamental and second harmonic waves. Such a picture would show that in any cycle of the first harmonic there are two cycles of the second, that is two crests and valleys, and that the crests of the second harmonic coincide with the zeros of the fundamental. The second harmonic generated in the first crystal is locked in this way to the phase of the incident fundamental, and then propagates through the intervening gas. On the other hand, the second harmonic generated in the second crystal is locked to the first harmonic entering that crystal, which has already propagated through the intervening gas. The contribution from optical propagation in the intervening gas to the phase of the output waves is generated at the first harmonic in one case and at the second harmonic in the other. Thus the output depends on the wavelength dispersion of the gas between the crystals. Careful theoretical calculations confirm this, and agree with the experimental results of

Summers and Velsco.

We conclude that for peak conversion the second harmonic waves generated in the two crystals must have the same phase as they exit the second crystal. This phase depends on both the dispersion of the intervening air, as well as the relative sign of the d-coefficients of the two crystals. It also depends on the phase mismatch, but this is too subtle to discuss in detail here.

The phase difference between the second harmonic and the fundamental from an air path of length L is

$$\delta \phi = \frac{4\pi L}{\lambda_0} (n_{2\omega} - n_{1\omega}) \quad (1)$$

For efficient conversion this phase must be either π or zero, depending on the relative sign of the effective d-coefficients in the two doubling crystals. If the relative sign is negative the phase change arising from the air must be π . For dry air at STP a phase change of π corresponds to a path length of about 7 cm. The relative sign of the d-coefficients can be determined from the standard expressions relating d_{eff} to d_{36} ; it depends on whether the XY-cuts of the crystals are the same, and on the relative orientation of the positive z-axes of the two crystals. The relative sign of the d-coefficients can be written

$$\frac{d_1}{d_2} = \xi_{xy} - \xi_Z \quad (2)$$

Here ξ_{xy} is -1 if the crystals have different XY-cuts, and +1 if they are the same. ξ_Z is +1 if the crystal Z-axes are parallel. The standard alternating-Z configuration is obtained from this parallel case by rotating one of the crystals about the beam direction. Then ξ_Z is +1 for the standard alternating-Z configuration also. If either z-axis is exactly reversed from one of these configurations an additional factor of -1 appears for type I crystals, +1 for type II crystals. Thus for same-cut crystals in the alternating-Z configuration the phase shift in the air gap must be 0 or 2π . If one Z-axis is reversed for type I crystals, the phase shift must be an odd multiple of π .

For the crystals that we tested, the required phase shift was 2π . For dry air at atmospheric pressure and at 24 °C, Velsco found that a 2π phase shift occurred for a change in crystal separation of 13.1 cm. Summers reported a value of 15 cm for laboratory air. The differences between dry and laboratory air are significant. In the experiment, we used the 13.1-cm spacing of the doublers from the more recent study[6].

VI. Example Calculations for the alternating-Z tripler.

This section contains graphs of conversion efficiency that demonstrate the principal characteristics of an alternating-Z arrangement with 80% deuterated crystals. The calculations were done with a plane-wave code for pulses with rectangular waveform. For all calculations, the doublers had thicknesses of 13 and 10 mm. Their angular offsets and physical separation were varied. The mixers were detuned by 60 μ rad and had thickness of either 10 or 13 mm. The conversion efficiencies depend partially on transmittances. We used single-surface transmittances that were deduced from measured transmittance curves for the 13-10-10 crystal set that was experimentally tested [7].

Figure 1 illustrates the general shape of the optimized conversion curve, and the variation of the intensity band with mixer thickness. Both calculations were done with a 13-10 doubler, with detunings of 400 and -520 μ rad. The mixers had thicknesses of 10 and 13 mm. It was assumed that proper correction for dispersion was provided by the air gap between the doublers. The dynamic range at 60% conversion is 9 for the 13-10-13 set and 10 for the 13-10-10 set. The peak efficiencies are about the same as that available from 2-crystal convertors. Therefore, the primary advantage is increased dynamic range.




Fig. 1. Conversion efficiency of optimized alternating-Z triplers.

The effect of dispersion in the air gap between the doublers is illustrated in Fig. 2. Calculations are for a 13-10-10 tripler with doubler offsets of +425 and -520 μrad , which were the conditions for the experimental test. At high intensity, a small phase error causes a large reduction in efficiency. Assuming that a spacing of about 13 cm provides a 2π phase shift [6], the errors of 0.15 and 0.20 radians correspond to spacing errors of 6.2 and 8.3 mm, respectively. Therefore, optimal usage of this arrangement in air would require careful spacing, but spacing should not be of concern if the tripler is in vacuum.




Fig. 2. Variation of conversion efficiency with phase lag of the 527-nm wave due to dispersion in the air gap between the doublers.

One of the most interesting characteristics of the alternating-Z arrangement is that best conversion, over the entire range of input design intensities, is obtained for unique detunings of the doublers. Small variations in doubler offsets cause variations of efficiency for the largest input intensities, see Fig. 3. The effect of severe departure from the optimum offsets are illustrated in Fig. 4. For the calculations in this figure, the offset of the 13-mm doubler was held at 400 μrad , and the offset of the 10-mm doubler was varied from 0 to -520 μrad .

The uniqueness of the detunings might be of value in a laser such as NIF. One would need a technique for correctly setting the orientations of the crystals, but this is true for all convertors. However, it would not be necessary to contemplate variation of the orientations when the waveform or intensity of the 1053-nm pulse was changed.

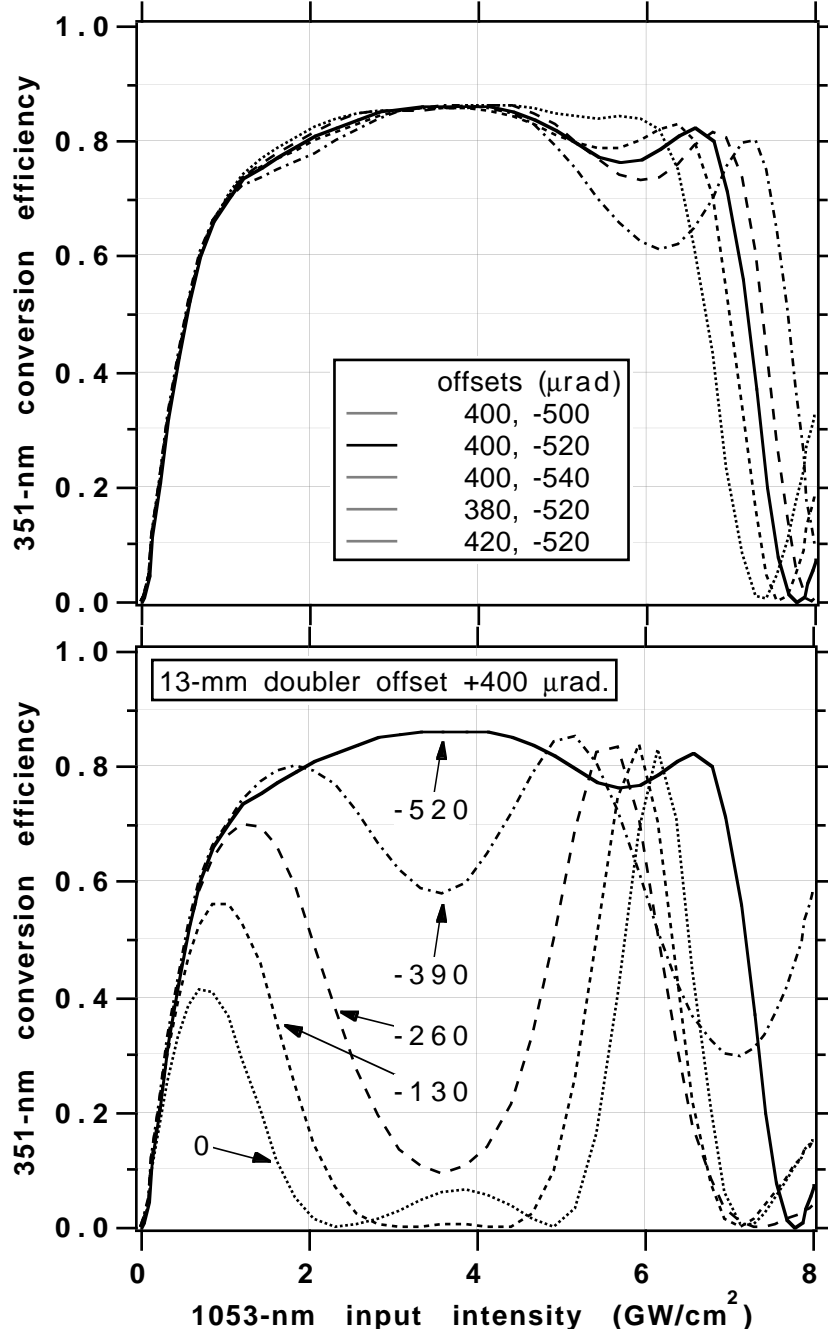


Fig. 3. Effect of small variations in the angular offsets of the doublers.

Fig. 4. Effect of large variations in offset angle of the 10-mm doubler.

VI. Experimental Arrangement and Measurements.

Figure. 5 shows a schematic diagram of the tripling experiment. The diagnostics package for the input beam recorded the fluence distribution in the beam, and the waveform and energy of the pulse. The 1053-nm input beam was 22 mm in diameter. A typical fluence distribution is shown in Fig. 6. The beam was shaped by apodization and image relaying. The local intensity modulation consisted of diffraction rings that were a residual of the beam apodization, and random structure that was caused by imperfections in optical components within the chain.




Fig. 5. Experimental arrangement.




Fig. 6. Fluence distribution in the input 1053-nm beam.

Waveforms were recorded by a Hamamatsu photodiode and a Tektronix SCD5000 digitizer. Experiments were done with both 1-ns and 6-ns pulses. Typical waveforms are shown in Fig. 7. Both of these waveforms were shaped by a Pockels cell shutter, and difficulties with the shutter caused a measurable shot-to-shot variation of the waveforms.




Fig. 7. Waveforms of the input 1053-nm pulses.

An absorbing-glass calorimeter was used to measure the fraction of the input energy that was contained in a central 1.3-cm-diameter sector of the input beam. This option provided a simple, spatially-averaged measurement of fluence in the center of the beam, and avoided the necessity of modeling the conversion of the varied intensities in the shoulders of the spatial distribution. A discussion of the uncertainties in fluence measurements is given in the appendix.

The output diagnostic package is also shown in Fig. 5. A sample of the output beam was reflected by the front surface of a bare silica wedge, and passed through an aperture with diameter of 1.3 cm. The apertured beam was dispersed by two silica prisms, and the energy and fluence distribution were measured for each of the harmonics. Spatially averaged output fluences were calculated from the area of the aperture and the energies.

The harmonic conversion crystals had aperture of 50 x 50 mm, and they were fabricated by diamond turning from 80% deuterated KD*P. The arrangement of the crystals is shown in Fig. 8. Each crystal was housed in a thermal enclosure and maintained at a temperature of 24.2 ± 0.1 °C, which was about 2 degrees above room temperature. Each oven was mounted in a gimbal that allowed angular tuning, and on a precision translation stage that allowed motion of the crystal in and out of the beam. The separation of the interior faces of the two doublers was 13.1 cm. This is the coherence length that Velsco and Eimerl measured in air at 24 °C and 760 Torr for the wavelengths of 1064 and 532 nm [6].

The crystals were oriented by rocking curves generated at high input intensities, about 3 GW/cm² for doublers and 2 GW/cm² for the tripler. The doublers were independently rocked, and the tripler was rocked while one of the aligned doublers was in the beam. The gimbals were used to set the detunings (internal) of +420 μ rad for the




Fig. 8. Arrangement of the convertor crystals

13-mm doubler and -520 μrad for the 10-mm doubler. The uncertainty of the angular alignment has three components, the pointing stability of the laser ($\pm 10 \mu\text{rad}$), the precision of the curve fit to the rocking data ($\pm 10 \mu\text{rad}$), and the accuracy of the detuning by the gimbal ($\pm 5 \mu\text{rad}$). Under the assumption that these are independent, the alignments should have been accurate to within $\pm 15 \mu\text{rad}$ during a typical shot.

Two experiments were done with 1-ns pulses. The tripler was translated out of the beam, and the mix ratio of the 1053- and 527-nm light was measured as a function of input intensity. The tripler was reinstalled, and conversion to 351-nm was measured with 1-ns input pulses that had peak intensities from 0.5 to 5.5 GW/cm^2 , and with 6-ns pulses that had intensities below 1 GW/cm^2 . Use of the longer pulses provided better accuracy in fluence measurements at low intensity.

VII. Analysis of the data

Two difficulties were encountered in the modeling. Because there were measurable differences in the waveforms of the pulses, neither fluence nor intensity was a pure variable. The data were reduced on a shot-by-shot basis using the 1053-nm waveform that was recorded for each shot. In graphs, the intensity for each shot was defined to be the average during the central one third of the waveform.

The second issue was dispersion in the air space between the two doublers. All of the data for high input intensity were inconsistent with results of calculations that used the measured tuning angles of the doublers, +420 and -520 μrad , and the assumption that dispersion in the air gap was eliminated by the spacing of 13.1 cm.

Assuming a phase lag of -0.15 radians for the 527-nm wave, at the input to the second doubler, provided a good overall description of the measured data.

The conversion data for the 6-ns pulses are shown in Fig 9. The agreement between measured efficiencies and those calculated from individual waveforms is excellent, and for these long pulses, the efficiencies are indistinguishable from those calculated with a rectangular waveform. The ability to model this data indicates that we were using the correct value of transmittances, thicknesses, and coupling coefficients in the model.




Fig. 9. Efficiency measured with low-intensity 6-ns input pulses.

The mix ratios that were measured with 1-ns pulses are shown in Fig 10. The mix ratio rose rather slowly with increase in input intensity, and had a positive slope when it passed through the optimum value of 2:1. The modeling indicates that the slow rise was caused by the large variation of intensity in the waveforms of the 1-ns pulses. The previously mentioned phase lag of the 527-nm beam adequately explains the high value of the mix ratio at high intensity. It was also possible to explain the high intensity values by assuming that the detunings of the doublers, +420 and -520 mrad, were in error by about 40 mrad. This angular error is larger than can be defended, and the assumption of a phase shift provided an explanation of all of the remainder of the data, so we believe it is the better of the options.




Fig. 10. Ratio of 527-nm to 1053-nm pulse energies that were measured with 1-ns pulses. Calculations assume a phase error of -0.15 radians in the air gap between the doublers, and detunings of +420 and -520 μ rad.

Figure 11 shows the conversion efficiencies that were measured with 1-ns pulses, and the efficiencies that were calculated from the 1053-nm calibrated waveforms. For an input pulse with peak intensity of about 3 GW/cm², the energy efficiency was 75% and the peak efficiency was 80%. Peak conversion of 83% was predicted by the code for rectangular pulses. For input pulses with peak intensities of 5-5.5 GW/cm², the largest




Fig. 11. 351-nm conversion efficiency for the 1-ns pulses. Calculations assume a phase error of -0.15 radians in the air gap between the doublers, and detunings of +420 and -520 μ rad.

conversion efficiency occurred at some time during the rise and fall of the pulse. The only significant discrepancy is one datum for input of 5.5 GW/cm².

Also shown in Fig. 11 is the conversion efficiency that was calculated under the assumption that all the waveforms were identical. One waveform was selected, and replicas were generated that had peak intensities ranging from 0.25 to 6 GW/cm². The conversion efficiencies calculated from these scaled waveforms are almost identical to those calculated from individual calibrated waveforms.

The conversion code calculated the fluence in all three harmonics, so it was possible to compare all of the measured fluences with the model. These results are shown in Fig. 12. The solid curves are fits through all of the data, either calculated or measured, for a particular harmonic. The agreement is excellent for the 351-nm wave, and acceptable for the residual 527-nm and 1053-nm light because the uncertainty is relatively large for calorimetry of low energy pulses.




Fig. 12. Output fluence at each of the harmonics. Measured and calculated results are indicated by, respectively, solid and open symbols. Curves were fitted through the data to ease distinction of the symbols.

VIII. Appendix. Uncertainty in fluence measurements.

All of the experimental data are either spatially averaged fluences, or ratios of these fluences. The uncertainties for the fluences were derived from a previous study of the precision of the calorimeters. We used four nominally identical absorbing-glass calorimeters (Scientech, Inc. Model 38-0111). They were electrically calibrated in a laboratory that is operated by the Nova laser staff. The electrical calibration was precise

to within less than 0.25%. The calibration is also highly accurate, provided that the exponentially decaying output signal is interpreted in the same manner during both calibration and usage

The calorimeter signals were amplified by Keithley Model 182 digital volt meters, and read by a 12-bit digitizer in LabView. The digitizer provided adequate resolution for signals with amplitude larger than about 0.25 volts, and the gain of the volt meter was adjusted to produce signals larger than this. Increasing the gain to allow measurements of small energies increases the sensitivity of the calorimeter to thermal noise, so measurements of small energies are inherently less precise.

The dependence of calorimetric precision on received energy was studied by the use of the arrangement that is shown in Fig 13. The arrangement used bare silica splitters, oriented at measured incidence angles, so the distribution of energy from the primary pulse into the two calorimeter channels could be accurately calculated. Both calorimeters measured the energy in the primary pulse, but one calorimeter received 30 times less energy.

Results of a test of two calorimeters are shown in Fig. 14. The ratio of the two measurements of primary-pulse energy was unity to within 0.25% when both calorimeters received at least 30 mJ. The value of unity for the ratio indicates consistency in the calibrations, and implies that any absolute uncertainty is a multiplicative factor that is common to both instruments.

Envelope curves were fitted to the data in Fig. 14, and used to assign an imprecision to each measurement of fluence in the conversion experiment. The imprecision is the primary factor in the uncertainty of ratios such as conversion efficiency. Systematic error is rigorously canceled, although it still appears as a small rescaling of the intensity or fluence axis. In contrast, estimates of systematic error have a strong effect on modeling, which must necessarily deal with the variation of absolute power with time in the calibrated waveform of the input pulse.




Fig. 13. Arrangement for using calorimeters that receive energies different by a factor of about 30 to measure the energy in the primary pulse.




Fig. 14. Ratios of the two measurements of pulse energy as a function of the energy that was recieved by the weakly illuminated

IX. Acknowledgements.

We are grateful for the assistance of Walter Sell who maintained and operated the OSL facility during these experiments, and recorded much of the data.

X. References.

1. W. Koechner, Solid-State Laser Engineering, Chapter 10, Springer-Verlag, 1988
2. D. Eimerl, "High average power harmonic conversion," IEEE J. Quant. Electron., vol. QE-23, No. 5, 575-593, May, 1987.
3. D. Eimerl, "Quadrature frequency conversion," IEEE J. Quant. Electron., vol. QE-23, No. 8, 1372-1378, Aug., 1987.
4. R. S. Craxton, "High efficiency frequency tripling schemes for high-power Nd:glass lasers," IEEE J. Quant. Electron. vol. QE-17, No. 9, 1771-1782, Sept. 1981.
5. M. A. Summers and R. D. Boyd, To our knowledge, this work was never published. The results are available in "Frequency conversion-experiments and theory, a collection of coversion studies for Nova."
6. S. P. Velsco and D. Eimerl, "Precise measurements of optical dispersion using a new interferometric technique," Appl. Opt., vol. 25, No. 8, 1344-1349, April 1986.
7. D. Milam, "Single-surface transmittance of AR-coated KDP crystals," Memo to distribution, LST-BLT96-076/jb, L-22593-01, July 1, 1996.